On the elicitation of indirect preferences in interactive evolutionary multiple objective optimization

Nominated for the best paper award

Michał Tomczyk, Miłosz Kadziński

Laboratory of Intelligent Decision Support Systems

Institute of Computing Science

Poznan University of Technology

July 8, 2020

Introduction

Preference vs. non-preference-based EMOAs

Observation: it is not practical to approximate an entire Pareto Front (PF) since the DM is interested in finding solutions that are relevant to him of her.

Preference-based EMOAs use the preference information provided by the DM to:

- **constrain** the search space, thereby **reducing the complexity of the problem;**
- impose an additional selection pressure, driving in this way the population towards **a region in the PF being highly preferred by the DM.**

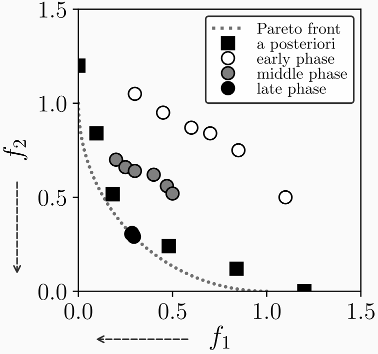

Considered class of preference-based EMOAs

Considered class of preference-based EMOAs

We consider the preference-based algorithms that have the following properties:

- **Accepted forms of the DM's preference:** indirect preference information, e.g., pairwise comparisons of solutions, class assignments, and so on
- **Preference elicitation policy:** interactive methods.
- **Exploitation of the DM's preferences:** robust techniques that exploit the whole space of compatible preference model instances to assess solutions consistently with the DM's preferences.

Reminder on IEMO/D

M. K. Tomczyk and M. Kadziński. Decomposition-based interactive evolutionary algorithm for multiple objective optimization. IEEE Transactions on Evolutionary Computation 24, 2 (2020), 320–334.

Reminder on IEMO/D

Characteristics of IEMO/D

- based on MOEA/D; hence, it implements the decomposition paradigm
- interactive algorithm
- uses L_α-norms to model the DM's value system
- incorporates the DM's preference information in the form of pairwise comparisons
- exploits the whole space of compatible model instances, i.e., L_α-norms, to derive robust recommendations

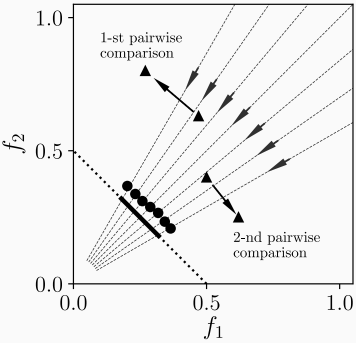

IEMO/D: Preference model & preference information

Preference model

$$L_\alpha^w(s) = \begin{cases} \left[\sum_{i=1,\dots,M} \left(w_i \dfrac{s_i - u_i}{n_i - u_i}\right)^\alpha\right]^{1/\alpha} & , \text{ for } \alpha < \infty, \\ \max_{i\in\{1,\dots,M\}} \left\{w_i \dfrac{s_i - u_i}{n_i - u_i}\right\} & , \text{ for } \alpha = \infty, \end{cases}$$

where w is a set of parameters, i.e., the *a priori* unknown objective weight vector.

The role of a preference preference model is to:

- **represent/model** the DM's value system
- **evaluate solutions** in the population consistently with the DM's preferences

IEMO/D: Preference model & preference information

Preference information

▶ The DM's pairwise comparisons $\mathcal{H} = \{(s^j \succ s^k)_1, \ldots, (s^j \succ s^k)_K$ can be used **to constrain the space of weight vectors** in the following way:

$$\bigvee_{s^j \succ s^k \in \mathcal{H}} L_\infty^w(s_j) < L_\infty^w(s_k).$$

▶ Any weight vector – and thus parameterized L_α-norm – that satisfies the above set of constraints is called **feasible or compatible**. Any compatible L_α-norm can be used during the evolutionary search to **reliably evaluate solutions** as such function potentially **represents the true DM's value system**.

▶ IEMO/D uses a representative subset of the space of compatible model instances to drive the evolutionary search towards the DM's **potentially the most relevant Pareto optimal solutions**.

IEMO/D: Intuitive example

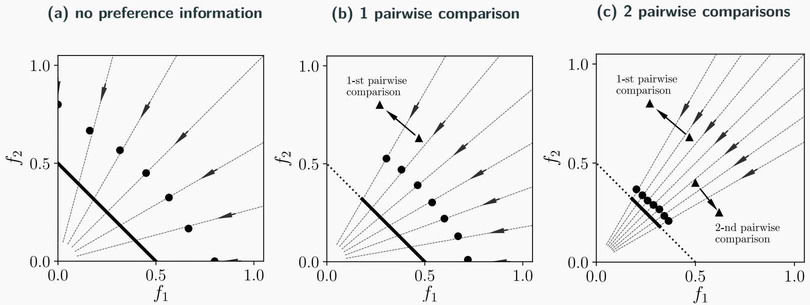

Research questions concerning preference learning in EMO

Research questions

1. How to select the reference solutions to be critically judged by the DM so as to maximize the information gain from his/her opinions?
2. How to distribute the interactions with the DM so as to make the best use of his/her feedback?
3. How can the DM be given more freedom in expressing his/her opinions?

IEMO/D★ is considered as a base for the research

★ *M. K. Tomczyk and M. Kadziński. Decomposition-based interactive evolutionary algorithm for multiple objective optimization. IEEE Transactions on Evolutionary Computation 24, 2 (2020), 320–334.*

Developments & results of experiments

1. Selection procedures

- **Random selection (RAND)**: solutions to be compared by the DM are selected randomly *(this is a benchmark procedure)*.
- **Medoid Selection (MS)**: Select two solutions representing two, vastly different regions in the objective space, using a K-medoids-inspired algorithm:

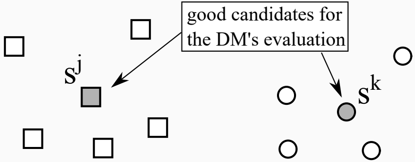

- **Pairwise Winning index based on Targets (PWIT)**: Selects two solutions so that the uncertainty of the DM's answer is the greatest. Let $PWI(s^j, s^k)$ be a share of compatible model instances maintained by IEMO/D that judge s^j better than s^k. Then, PWIT selects solutions as follows:

$$(s^j, s^k) \leftarrow \text{argmax}_{s^j, s^k \in P} min\{PWI(s^j, s^k), PWI(s^k, s^j)\}$$

Brief note on experimental setting

- the performance of IEMO/D was verified on WFG benchmarks
- the number of objectives ranged from 2 to 5
- the algorithm was run multiple times to obtain reliable results
- L_α-norm was used to simulated the DM's answers on indirect questions
- the same function was also employed as an "oracle", i.e., it was used to assess solutions in the the population; specifically, two performance indicators were used
    - BRSD: relative difference between the DM-perceived quality of the best solution in the population and the best solution in the PF.
    - ARSD: average relative difference between the DM-perceived quality of constructed solutions and the best solution in the PF.

Example results

- average results attained by IEMO/D applied to WFG3 with 3 objectives
- the interactions with the DM were performed 12 times, at regular intervals

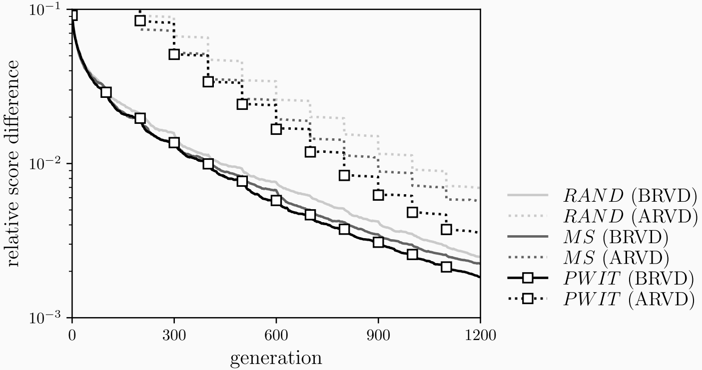

- the best results were attained by IEMO/D coupled with **Pairwise Winning Index based on Targets (PWIT)** selection procedure

2. Interaction patterns

- **Regular intervals (RI)**: the interactions with the DM are distributed evenly throughout the evolutionary search *(this is a benchmark procedure)*.
- **Improved Goals ($\boldsymbol{IG^{th}}$)**: the interactions are performed when **the efficiency of generating offspring that yields an improvement is particularly low**. In this way, the additional selection pressure can be imposed to enhance the search process. To decide upon the questioning of the DM, **the share of optimization goals that have been updated from one generation** to another is compared with some fixed threshold th. If the computed share **decreases below $\boldsymbol{th}$**, the interaction is triggered.

2. Interaction patterns

- **Regular intervals (RI)**: the interactions with the DM are distributed evenly throughout the evolutionary search *(this is a benchmark procedure)*.
- ~~**Improved Goals (IG^{th})**~~
- **Geometric Progression based on Improved Goals(GP-IG^{th})**: Uses the (IG^{th}) to trigger only the first interaction. The remaining interactions are distributed so that thus formed elicitation intervals are determined by a geometric progression.

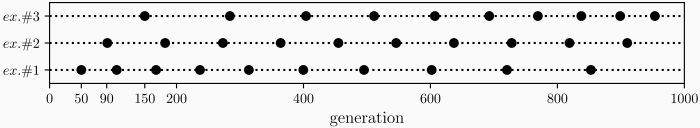

Example distributions imposed by GP-IG^{th} procedure

Example results

- the average number of interactions performed by IEMO/D when applied to WFG3 with 2 objectives
- the limit for the number of interactions with the DM was set to 12.

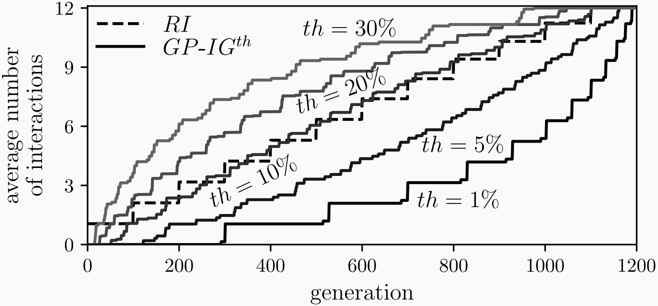

- the interactions are postponed when the threshold *th* is set to a low value; when compared to the Regular Intervals (RI) interaction pattern...
- ... and vice versa.

Example results

- the average ARSD attained by IEMO/D when applied to WFG3 with 2 objectives
- the limit for the number of interactions with the DM was set to 12

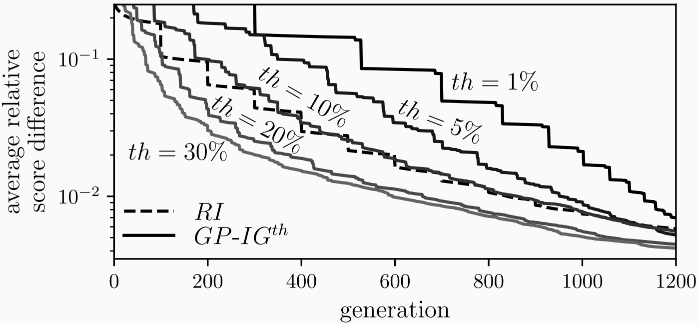

- The best results were attained by IEMO/D that performed interactions with the DM more frequently **at the early stage of the optimization process**.

3. Accepted forms of indirect DM's preferences

IEMO/D was revisited to accept different forms of indirect preference information:

- **Pairwise comparison:** $s^j \succ s^k$
- **Complete order of $\boldsymbol{k} = 3$ solutions:** $s^j \succ s^k \succ s^l$
- **Selection of the most preferred option in $\boldsymbol{k} = 3$ solutions:** $s^j \succ s^k \wedge s^j \succ s^l$
- **Preference intensity:** $s^j \succ^{Cl_l} s^k$, $l \in \{1 = \text{weak}, 2 = \text{medium}, 3 = \text{strong}\}$

The impact of incorporating different indirect preferences on the search space reduction – example results

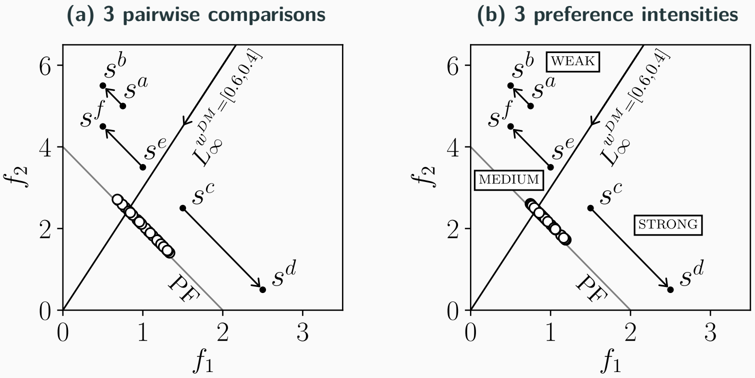

3. Accepted forms of indirect DM's preferences

The impact of incorporating different indirect preferences on the search space reduction – example results

(a) best individual in 3 solutions **(b) a complete order of 3 solutions**

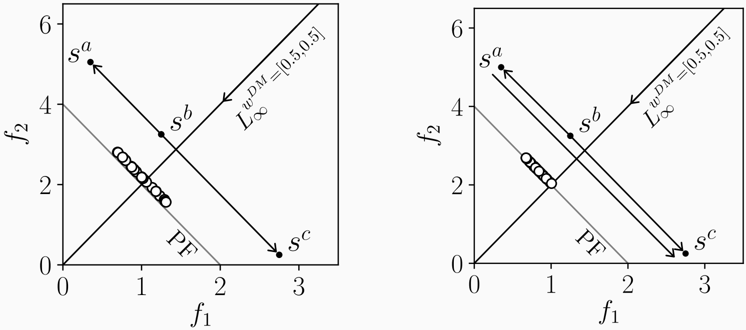

3. Accepted forms of indirect DM's preferences

The impact of incorporating different indirect preferences on the DM-perceived quality of constructed solutions

- **Pairwise comparisons:** $s^j \succ s^k$
  Cognitive cost $= 1$
- **Complete order of $k = 3$ solutions:** $s^j \succ s^k \succ s^l$
  Cognitive cost $= 3$
- **Selections of the most preferred option in $k = 3$ solutions:** $s^j \succ s^k \wedge s^j \succ s^l$
  Cognitive cost $= 2$
- **Preference intensities:** $s^j \succ^{Cl_l} s^k$, $l \in \{1 = \text{weak}, 2 = \text{medium}, 3 = \text{strong}\}$
  Cognitive cost $= 1$

Experimental setting

- the limit for the total cognitive effort was set to 12
- the interactions were distributed evenly

Example results

- average results attained by IEMO/D when applied to WFG3 with 3 objectives.

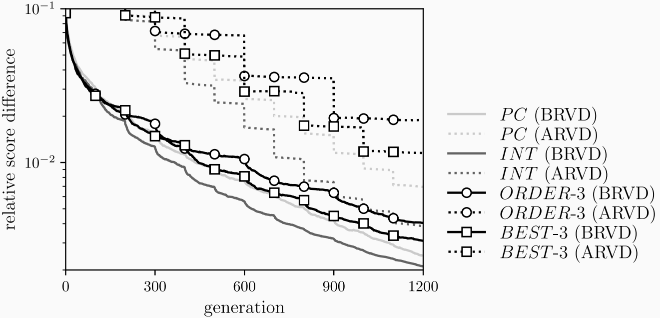

- the worst results were attained by IEMO/D supplied with **the complete orders (*ORDER*-3) or the selections of the best individuals (*BEST*-3)**
- in turn, IEMO/D coupled with the less informative but more frequently provided **pairwise comparisons (*PC*) and preference intensities (*INT*) attained the best results**.

Summary

Summary

This study concerned interactive EMO implementing the paradigm of active preference learning and incorporating robustness concern. In this work, three essential challenges related to the elicitation of indirect preference information in EMO were addressed:

- **Selecting reference solutions to be compared by the DM**
  We showed that the performance of the algorithm may be improved when the solutions to be compared by the DM are selected so as to maximize information gain from his/her answer.
- **Deciding when he/she should be questioned**
  We showed that the performance of the algorithm may be improved when the interactions are suitably adjusted according to properties of the underlying problem.
- **Accepting different forms of indirect preferences** We compared the performance of IEMO/D when supplied with different indirect preferences. The results revealed that it is more beneficial to ask for less demanding preferences more frequently, than vice versa.

References

- **IEMO/D algorithm**
  M. K. Tomczyk and M. Kadziński. Decomposition-based interactive evolutionaryalgorithm for multiple objective optimization. IEEE Transactions on Evolutionary Computation 24, 2 (2020), 320–334.
- **Preference elicitation strategies**
  M. Kadziński, M. K. Tomczyk, and R. Słowiński. Preference-based cone contraction algorithms for interactive evolutionary multiple objective optimization. Swarm and Evolutionary Computation 52 (2020), 100602.
- **Preference elicitation strategies & inconsistency analysis**
  J. Marquis, E. S. Gel, J. W. Fowler, M. Köksalan, P. Korhonen, and J. Wallenius. Impact of Number of Interactions, Different Interaction Patterns, and Human Inconsistencies on Some Hybrid Evolutionary Multiobjective Optimization Algorithms. Decision Sciences 46, 5 (2015), 981–1006.
- **A study on inconsistency analysis &**
  techniques for selecting solutions to be compared by the DM
  M. K. Tomczyk and M. Kadziński. EMOSOR: Evolutionary multiple objective optimization guided by interactive stochastic ordinal regression. Computers & Operations Research 108 (2019), 134–154.

Thank you for your attention ☺
Do you have any questions?

michal.tomczyk@cs.put.poznan.pl